

Caching and Auditing in the RPPM Model

Jason Crampton and James Sellwood

Information Security Group
Royal Holloway, University of London

December 1, 2014

Abstract

Crampton and Sellwood recently introduced a variant of relationship-based access control based on the concepts of relationships, paths and principal matching, to which we will refer as the RPPM model. In this paper, we show that the RPPM model can be extended to provide support for caching of authorization decisions and enforcement of separation of duty policies. We show that these extensions are natural and powerful. Indeed, caching provides far greater advantages in RPPM than it does in most other access control models and we are able to support a wide range of separation of duty policies.

1 Introduction

Whilst the majority of computer systems employ some form of role-based access control model, social networking sites have made use of the relationships between individuals as a means of determining access to resources. Recent work on relationship-based access control models has attempted to further develop this concept but has frequently remained focused on the relationships that exist between individuals [4, 9]. Crampton and Sellwood define a more general model for access control utilising relationships between entities, where those entities can represent any physical or logical component of a system [7]. These entities and their (inter-)relationships are described by a multigraph, called the *system graph*. Authorization requests in the RPPM model are processed by first determining a list of matching principals. This list of principals is identified using principal-matching rules and the system graph. Once a list of matched principals is determined, the specific action in the request is authorized or denied based on authorization rules defined for those principals and the object.

The RPPM model provides the necessary foundations for general purpose relationship-based access control systems, but there are a number of simple enhancements which would greatly increase its utility and efficiency. The evaluation of path conditions can be complex in system graphs containing many nodes of high degree. Support for caching of previously matched principals would significantly reduce the processing necessary during the evaluation of an authorization request. The introduction of caching support is, therefore, our first enhancement.

Our second enhancement adds support for request evaluation audit records to be kept, such that future authorization requests may be evaluated both on the current relationships within the system graph but also using historical information about past actions by subjects. Such mechanisms allow us to support constraints such as separation of duty and Chinese Wall policies and lay a foundation for future work on workflow authorization using the model.

The rest of this paper is arranged with background information on the RPPM model provided in Section 2 and then the two enhancements described individually in Sections 3 (caching) and 4 (audit records). We discuss related work in Section 5 and draw conclusions of our contributions and identify future work in Section 6.

2 The RPPM Model

The RPPM model, described in detail in [7], employs a system graph to capture the entities of a system and their (inter-)relationships. The entities (physical or logical system components) are nodes within the system graph whilst the relationships are labelled edges. The system graph’s ‘shape’ is constrained by a system model, which identifies the types of entities and relationship which are supported. It does so by defining a permissible relationship graph whose nodes are the possible types of entities in the system graph and whose labelled edges indicate the relationships which *may* exist in the system graph between entities of the connected types.

Definition 1. *A system model comprises a set of types T , a set of relationship labels R , a set of symmetric relationship labels $S \subseteq R$ and a permissible relationship graph $G_{PR} = (V_{PR}, E_{PR})$, where $V_{PR} = T$ and $E_{PR} \subseteq T \times T \times R$.*

Definition 2. *Given a system model (T, R, S, G_{PR}) , a system instance is defined by a system graph $G = (V, E)$ where V is the set of entities and $E \subseteq V \times V \times R$. Making use of a mapping function $\tau : V \rightarrow T$ which maps an entity to its type, we say G is well-formed if for each entity v in V , $\tau(v) \in T$, and for every edge $(v, v', r) \in E$, $(\tau(v), \tau(v'), r) \in E_{PR}$.*

Within the RPPM model, authorization requests have the form $q = (s, o, a)$, where a subject s requests authorization to perform action a on target object o . The authorization policy is abstracted away from subjects by the use of security principals. These principals are matched to requests through the satisfaction of path conditions using edges in the system graph, where a path condition π represents a sequence of relations with specific labels from the set R .

Definition 3. *Given a set of relationships R , we define a path condition recursively:*

- \diamond is a path condition;
- r is a path condition, for all $r \in R$;
- if π and π' are path conditions, then $\pi ; \pi'$, π^+ and $\bar{\pi}$ are path conditions.

A path condition of the form r or \bar{r} , where $r \in R$, is said to be an edge condition.

Informally, $\pi ; \pi'$ represents the concatenation of two path conditions; π^+ represents one or more occurrences, in sequence, of π ; and $\bar{\pi}$ represents π reversed; \diamond defines an “empty” path condition.

Definition 4. *Given a set of relationships R , we define a simple path condition recursively:*

- \diamond , r and \bar{r} , where $r \in R$, are simple path conditions;
- if $\pi \neq \diamond$ and $\pi' \neq \diamond$ are simple path conditions, then $\pi ; \pi'$ and π^+ are simple path conditions.

A path condition can describe highly complex and variable-length paths within the system graph. However, Crampton and Sellwood proved that every path condition can be reduced to an equivalent simple path condition [7, §2.2], thereby simplifying the design of the principal-matching algorithm.

Definition 5. Given a system graph $G = (V, E)$ and $u, v \in V$, we write $G, u, v \models \pi$ to denote that G, u and v satisfy path condition π . Formally, for all G, u, v, π, π' :

- $G, u, v \models \diamond$ iff $v = u$;
- $G, u, v \models r$ iff $(u, v, r) \in E$;
- $G, u, v \models \pi ; \pi'$ iff there exists $w \in V$ such that $G, u, w \models \pi$ and $G, w, v \models \pi'$;
- $G, u, v \models \pi^+$ iff $G, u, v \models \pi$ or $G, u, v \models \pi ; \pi^+$;
- $G, u, v \models \bar{\pi}$ iff $G, v, u \models \pi$.

Definition 6. Let P be a set of authorization principals. A principal-matching rule is a pair (π, p) , where π is a path condition and $p \in P$ is the associated principal. A list of principal-matching rules is a principal-matching policy.

In the context of a principal-matching rule, a path condition is called the *principal-matching condition*.

The request and system graph are evaluated against the principal-matching policy utilising a *principal-matching strategy* (PMS) to determine the list of matched principals for the request. The PMS specifies how principal-matching rules should be evaluated, for example whether the first matching principal applies (in the case of the **FirstMatch** PMS) or whether all matching principals apply (**AllMatch**). A *default* principal-matching rule (\top, p') may, optionally, be employed as the last rule in the policy and will automatically result in its principal p' being matched whenever the rule is evaluated.

A system graph G , two nodes u and v in G , a principal-matching policy ρ , and a principal-matching strategy σ determines a list of principals MP associated with the pair (u, v) . We evaluate each principal-matching rule (π, p) in turn and add p to the list of matched principals if and only if $G, u, v \models \pi$. We then apply the principal-matching strategy to the list of matched principals to obtain MP . (Obviously, optimizations are possible for certain principal-matching strategies.) We write $G, u, v \xrightarrow{\rho, \sigma} MP$ to denote this computation.

Once determined, the list of matched principals is used to identify relevant authorization rules in the authorization policy.

Definition 7. An authorization rule has the form (p, o, a, b) , where p is a principal, o is an object, a is an action and $b \in \{0, 1\}$, where $b = 0$ denies the action and $b = 1$ grants the action. In order to ease authorization policy specification we allow for the use of \star instead of o , to represent all objects, or instead of a , to represent all actions. These global authorization rules, therefore, have the form (p, \star, a, b) , (p, o, \star, b) or (p, \star, \star, b) . An authorization policy is a list of authorization rules.

The matching of principals to authorization rules yields a list of authorization decisions, which is reduced to be single decision using a *conflict resolution strategy* (CRS). The CRS is used in much the same way as a rule-combining or policy-combining algorithm is used in XACML. It may specify that particular outcomes are prioritised, such as (**AllowOverride** or **DenyOverride**), or that the first conclusive decision should be used (**FirstMatch**).

To summarise, given a request (s, o, a) , where s and o are nodes in the system graph and a is an action, we first compute the list of matched principals $G, s, o \xrightarrow{\rho, \sigma} MP$. We then use MP and the authorization policy to determine which actions are granted and denied for those principals and apply the CRS to determine a final decision. In this paper, we assume the use of the **AllMatch** PMS and **DenyOverride** CRS throughout.

3 Caching

The most complex part of evaluating an authorization request in the RPPM model is the principal matching stage [7, §3]. This process attempts to satisfy path conditions within principal-matching rules using paths between the subject and the object of the request. It is important to note that the requested action is immaterial during this processing stage (only becoming relevant when the authorization rules are considered). The list of matched principals for a subject-object pair remains static until a change is made to the system graph or certain associated policy components. Even then, not all of the possible changes would impact the matched principals between a particular subject and object.

We introduce the concept of *caching edges* and make use of the relative stability of matched principals in order to reduce the processing required for future authorization requests. We first redefine the system graph to support these new edges. In particular, when we evaluate a request (s, o, a) that results in a list of matched principals MP , we add an edge (s, o, MP) to the system graph, directed from s to o and labelled with MP .

Informally, a caching edge (s, o, MP) directly links s to o and identifies the matching principals MP relevant to requests of the form (s, o, a) . The processing of subsequent authorization requests can skip the principal matching stage and use MP in conjunction with the authorization rules to evaluate a request of the form (s, o, a) .

To illustrate, consider the simple system graph G_1 , shown in Figure 1a, and the following principal-matching and authorization policies

$$\begin{aligned}\rho &= [(r_1, p_1), (r_2, p_2), (r_3, p_3), (r_1 ; r_3, p_4), (r_2 ; r_3, p_5)] \\ PA &= [(p_5, \star, a_1, 1), (p_5, \star, a_2, 0)].\end{aligned}$$

If an authorization request $q_1 = (v_2, v_4, a_1)$ is made, then $G_1, v_2, v_4 \xrightarrow{\rho, \sigma} [p_5]$, because the only principal-matching condition from the policy which can be satisfied between v_2 and v_4 in G_1 is $r_2 ; r_3$. Then the authorization rule $(p_5, \star, a_1, 1)$ applies and the set of possible decisions $PD = \{1\}$; thus the request is authorized. At this stage we may add a caching edge $(v_2, v_4, [p_5])$ to produce the system graph shown in Figure 1b. We use the convention that caching edges have a diamond-shaped arrow head.




Figure 1: Adding a caching edge

If an authorization request $q_2 = (v_2, v_4, a_2)$ is subsequently made, the caching edge $(v_2, v_4, [p_5])$ allows us to evaluate the request without re-evaluating the principal-matching policy; the authorization rule $(p_5, \star, a_2, 0)$ subsequently results in request q_2 being denied ($PD = \{0\}$).

To consider the scale of the potential benefit of caching edges, we review the experimental data reported by Crampton and Sellwood for numbers of nodes visited (n) and edges considered (e) during sample request evaluations (see Table 1 and [7, §3.3]). With support for caching edges, if the subject-object pairs participating in any of these requests were to be involved in subsequent requests the processing would instead be limited to locating the appropriate caching

edge. It should be clear that when considering requests which may require upwards of 50 edge evaluations (in a small example system graph), replacing this with a single caching edge lookup could dramatically improve evaluation performance.

Table 1: Experiment results from [7, §3.3]

Path condition	Request	n	e	Match Found
π_1	q_1	5	19	Yes
π_1	q_2	7	24	Yes
π_2	q_3	4	15	Yes
π_3	q_4	17	58	Yes
π_3	q_5	7	24	No

In the worst case, the number of caching edges directed out of a node is $O(|V|)$, where V is the set of nodes in the system graph. However, there are strategies that can be used to both prevent the system graph realizing the worst case and to reduce the impact of large numbers of caching edges. To maintain an acceptable number of caching edges, we could, for example, use some form of cache purging. We can also distinguish between relationship edges and caching edges using some flag on the edges and index the caching edges to dramatically decrease the time taken to search the set of caching edges. Employing these techniques should enable the benefits of caching edges to be realised without incurring unacceptable costs during identification of the relevant caching edge. Further experimental work is required to determine how best to make use of caching edges.

3.1 Preemptive Caching

Any optimisation provided by the caching of matched principals relies upon the existence of a caching edge in order to reduce the authorization request processing; the first request between a subject and object must, therefore, be processed normally in order to determine the list of matched principals which will label the caching edge. If this initial evaluation were only performed when an authorization request were submitted, then the benefit of caching edges would be limited to repeated subject-object interactions alone.

However, many authorization systems will experience periods of time when no authorization requests are being evaluated. The nature of many computing tasks is such that authorization is required sporadically amongst longer periods of computation by clients of the authorization system and idle time for the authorization system itself. These periods of reduced load on the authorization system can be employed for the purpose of *preemptive caching*.

Thus for pairs of nodes (u, v) in the system graph, we may compute $G, u, v \xrightarrow{\rho, \sigma} MP$ and insert a caching edge (u, v, MP) . The fact that a request's action is not employed during the principal matching process means that to perform this further optimization an authorization system solely requires a subject and object between whom the matched principals are to be identified. There are numerous potential strategies for determining which subject-object pairs should be considered for preemptive caching. Here we describe two simple and natural strategies.

Subject-focused. Subject-focused preemptive caching assumes that subjects who have recently made authorization requests are *active* and so will likely make further requests. The authorization system, therefore, prioritises determining the list of matched principals between the most recently active subjects and a set of target objects. The set of target objects could

be selected at random or may be systematically chosen using an appropriate mechanism for the system defined in the system graph. This might involve the target objects being *popular*, *significant* or those whose access may be particularly *time-sensitive*. We envisage that the interpretation of these concepts may be system specific, as may their relative worth.

As preemptive caching builds the number of caching edges within the system graph the number of subjects and objects under consideration could be expanded to provide greater coverage of the potential future requests.

Object-focused. In certain applications, there will be resources that will be used by most users, such as certain database tables. Thus, it may make sense to construct caching edges for all active users for certain resources.

No matter the strategy, preemptive caching makes use of available processing time in order to perform the most complex part of authorization request evaluation: principal matching. Any requests that are made utilising a subject-object pair which have already been evaluated by preemptive caching will be able to make use of the caching edge already established, even if that request were the first received for that pair. Once determined, caching edges resulting from preemptive caching are no different from those established as a result of request evaluation.

3.2 Cache Management

A change to any of the following components of the model could modify the list of matched principals for a subject and object:

- the system graph;
- the principal-matching policy;
- the principal-matching strategy.

Such changes, therefore, may affect the correctness of caching edges. (The obvious exception is a change to the system graph resulting from the addition or deletion of a caching edge.) The most crude management technique for handling such changes involves removing all caching edges from the system graph whenever one of the above changes occurs.

In certain specific scenarios it may be possible for a system to identify a scope of impact for a particular change and thus apply a more refined management technique. For example, if a change to the principal-matching policy removes all rules which are used to match a certain principal (and nothing more), then it would be sufficient for only caching edges labelled with a list including that principal to be purged. Whilst such a refinement may further optimise the operations performed by the authorization system, its applicability will depend upon the configuration of the authorization system in its entirety.

We have already noted that it may make sense to purge the cache in order to limit the number of caching edges in the system graph. Again, there are several possible purging strategies. One would be simply to set a maximum threshold for the number of caching edges in the system graph. A second, perhaps more useful, strategy would be to set a maximum threshold for the out-degree (measured in terms of caching edges) for any node in the graph. We may also “retire” caching edges: any edge that hasn’t been used as part of a request evaluation for some time period will be purged. And we could employ mixed strategies, which might depend on the application and the nature of the system graph.

4 Audit Records

Currently, the RPPM model’s authorization request processing is “memoryless” with respect to previous requests and their respective outcomes. Various scenarios and security policy principles make use of historical data. Reputation systems and history-based access control (HBAC) systems [1, 8, 14], for example, rely on knowledge of previous interactions and requests in order to correctly make authorization decisions. The Chinese Wall [3] and non-static separation of duty principles [11, 15] also rely on knowledge of previous actions to enforce their constraints.

We introduce the concept of *audit edges*, through which we track the outcomes of authorization requests for subsequent use in policy evaluation. Audit edges come in two flavours: those which directly record the decision of a previous authorization request (authorized and denied *decision audit edges*) and those which, more generally, record an entity’s interest in other entities based on its authorized requests (active and blocked *interest audit edges*). It should be noted that whilst we make direct use of audit edges for policy evaluation, they also have value in a system purely as an audit record. We extend the set of relationships and further redefine the system graph to support these new edges. Specifically, in the case of decision audit edges:

- for each action a , we define two relationships a^{\oplus} and a^{\ominus} and include the sets $\{a^{\oplus} : a \in A\}$ and $\{a^{\ominus} : a \in A\}$ in the set of relationships;
- if the decision for request (u, o, a) is allow, then we add the edge (u, o, a^{\oplus}) into the system graph;
- if the decision for request (u, o, a) is deny, then we add the edge (u, o, a^{\ominus}) into the system graph.

Both authorized and denied decision audit edges are inserted, automatically, into the system graph after request evaluation completes. If such an edge does not already exist, a decision audit edge is added between the subject and object of the evaluated request, indicating its result.

The addition of interest audit edges also occurs automatically after request evaluation completes. For such edges, the subject is the source node of the interest edge (as for decision audit edges); however, the destination node may not be the object of the request. Interest audit edges are discussed in more detail in Section 4.2.

4.1 Enforcing Separation of Duty

Separation of duty requires that certain combinations of actions are performed by a number of distinct individuals so as to reduce the likelihood of abuse of a system. In its simplest form, separation of duty constraints require two individuals to each perform one of a pair of distinct actions so that a single individual cannot abuse the system. A common application environment for such constraints is that of a finance system, where, for example, the individual authorized to add new suppliers should not be the same individual who is authorized to approve the payment of invoices to suppliers. If a single individual were able to perform both of these actions they could set themselves up as a supplier within the finance system and then approve for payment any invoices they submitted as that supplier. We define a mechanism here through which n individuals can be required to perform n actions on an object. Before doing so, we explain a simplified version of the mechanism for the case $n = 3$.

Let us consider the system graph G_2 (see Figure 2a), the principal-matching policy $\rho = [(r, p)]$ and the authorization policy $PA = [(p, o, \star, 1)]$. With these policies and without audit edges, if individual u_1 makes the request $q_1 = (u_1, o, a_1)$ this will be authorized by matching principal




Figure 2: Adding decision audit edges

p , as will subsequent requests $q_2 = (u_1, o, a_2)$ and $q_3 = (u_1, o, a_3)$. A similar result would have occurred if these requests had been submitted with u_2 or u_3 as the subject.

A basic implementation of separation of duty can be employed by introducing a new principal p_{seen} which matches if a user has performed any action on the object. We change the principal-matching and authorization policies to

$$[(a_1^{\oplus}, p_{seen}), (a_2^{\oplus}, p_{seen}), (a_3^{\oplus}, p_{seen}), (r, p)] \quad \text{and} \quad [(p_{seen}, o, \star, 0), (p, o, \star, 1)]$$

respectively¹. Using this combination of policies means that any user who has performed an action on object o is prevented from performing another action as all actions are denied to the principal p_{seen} .

Whilst this basic implementation fulfils the requirement that no user may perform more than one action on the object, we may wish to specify more refined separation of duty policies within the system. The basic implementation has several limitations which RPPM's policies are flexible and powerful enough to resolve. Specifically, all actions within the system are included in the separation of duty constraint due to the use of an authorization rule for all actions $(p_{seen}, o, \star, 0)$. Additionally, having performed an action on o a user is unable to repeat the action performed, as well as being unable to perform any other action.

If we wish to enforce a more flexible separation of duty constraint on a subset of actions $\{a_1, a_2, a_3\} \subseteq A$ such that distinct individuals are required to perform each action, we can modify the principal-matching policy to $\rho = [(a_1^{\oplus}, p_1), (a_2^{\oplus}, p_2), (a_3^{\oplus}, p_3), (r, p)]$ and the authorization policy to:

$$PA = [(p_1, o, a_2, 0), (p_1, o, a_3, 0), (p_2, o, a_1, 0), (p_2, o, a_3, 0), \\ (p_3, o, a_1, 0), (p_3, o, a_2, 0), (p, o, \star, 1)]$$

The first action. Revisiting our example for G_2 , an initial request $q_1 = (u_1, o, a_1)$ will, once again, be authorized (with $MP = [p]$) but will, this time, result in the addition of an authorized decision audit edge (u_1, o, a_1^{\oplus}) . If u_1 then makes a request $q_2 = (u_1, o, a_2)$ this will be denied as $MP = [p_1, p]$ and the authorization rule $(p_1, o, a_2, 0)$ indicates a deny which overrides the authorization from the rule $(p, o, \star, 1)$. Similarly if u_1 makes a request $q_3 = (u_1, o, a_3)$ this will be denied as once again $MP = [p_1, p]$ and the deny authorization rule $(p_1, o, a_3, 0)$ overrides $(p, o, \star, 1)$. These two denied requests would result in denied decision audit edges (u_1, o, a_2^{\ominus}) and (u_1, o, a_3^{\ominus}) .

The second action. However, if u_3 makes the request $q_4 = (u_3, o, a_2)$ this will be authorized with $MP = [p]$ and use of the authorization rule $(p, o, \star, 1)$; the authorized decision audit

¹We assume the use of the AllMatch PMS and the DenyOverride CRS, but we could equally employ the FirstMatch PMS with any CRS as long as we ensure that the constraint principal-matching rules are added before any existing rules.

edge (u_3, o, a_2^\oplus) results. If u_3 attempts to then make request $q_5 = (u_3, o, a_3)$ this will be denied in the same manner that request q_3 was, with the subsequent addition of a denied decision audit edge (u_3, o, a_3^\ominus) .

The last action. As a_1 was performed by u_1 and a_2 was performed by u_3 it remains, for successful operation, for u_2 to make request $q_6 = (u_2, o, a_3)$. This request will be authorized with $MP = [p]$ and the use of the authorization rule $(p, o, \star, 1)$, resulting in the authorized decision audit edge (u_2, o, a_3^\oplus) . The system graph that results after all of these requests have been made is as shown in Figure 2b.

More generally, suppose we have a principal-matching policy ρ and an authorization policy PA . If we require that the actions $\{a_1, \dots, a_n\}$ should each be performed by different users (and the same action may be repeated), we add the rules

$$(a_1^\oplus, p_1), \dots, (a_n^\oplus, p_n)$$

to ρ and let the new policy be ρ' . And for each principal p_i , we add the set of rules

$$\{(p_i, o, a_j, 0) : 1 \leq j \leq n, j \neq i\}.$$

to PA denoting the new policy PA' . We then have the following result

Proposition 1. *Given an RPPM separation of duty policy, as described above, for any user u the request (u, o, a) is allowed if the request is authorized by ρ' and PA' and no request of the form (u, o, a') , where $a' \neq a$, has been previously authorized; the request is denied otherwise.*

Proof. The proof proceeds by induction on the number of evaluated requests. Consider the (base) case when no requests have yet been made. A request (u, o, a) where $a \in \{a_1, \dots, a_n\}$ will not match any of the n inserted principal-matching rules as no decision audit edges currently exist in the system graph. Thus request (u, o, a) will be authorized if it is authorized by ρ and PA (and hence will be authorized by ρ' and PA').

Now suppose the result holds for all sequences of m requests and consider the request (u, o, a) where $a \in \{a_1, \dots, a_n\}$.

- If u has previously performed a constrained action a_i , $1 \leq i \leq n$, then the request will satisfy principal-matching condition (a_i^\oplus, p_i) .

Now, if $a_i = a$, there is no authorization rule of the form $(p_i, o, a_i, 0)$ and the request will, therefore, be authorized if and only if it is authorized by ρ and PA .

Conversely, if $a_i \neq a$, then $a = a_j$, for some $j \neq i$, and the authorization rule $(p_i, o, a, 0)$, together with the DenyOverride CRS will cause the request to be denied.

- If user u has not previously performed a constrained action then the request will not match any of the principal-matching rules that were added to create ρ' . Thus the request will only be authorized if it is authorized by ρ and PA .

□

4.2 Enforcing Chinese Walls

The Chinese Wall principle may be used to control access to information in order to prevent any conflicts of interest arising. The standard use case concerns a consultancy that provides services

to multiple clients, some of whom are competitors. It is important that a consultant does not access documents of company A if she has previously accessed documents of a competitor of A .

To support the Chinese Wall policy, systems classify data using conflict of interest classes [3], indicating groups of competitor entities. Requests to access a company's resources within a conflict of interest class will only be authorized if no previous request was authorized accessing resources from another company in that conflict of interest class.

Unlike the general approach for separation of duty, a general approach for Chinese Wall requires fewer policy changes but does rely on a particular basic layout of system graph. This layout is such that the users who will be making requests are connected (directly or indirectly) to the companies (which may or may not be competitors of each other). These companies are then connected to their respective data entities, which will be the targets of users' requests. This arrangement is depicted, conceptually, in Figure 3a, with the path condition π_1 representing the chain of relationships between users and companies and π_2 between the data entities and the companies.² In other words, the path from an authorized user to a company will contain the same labels (and will match the path condition π_1), irrespective of the specific identities of the user and company. Similarly, the path from a data object to its owner company will contain the same labels (and match the path condition π_2). Thus, the principal that is authorized to access companies' data objects would be matched using the path condition $\pi_1 ; \overline{\pi_2}$.

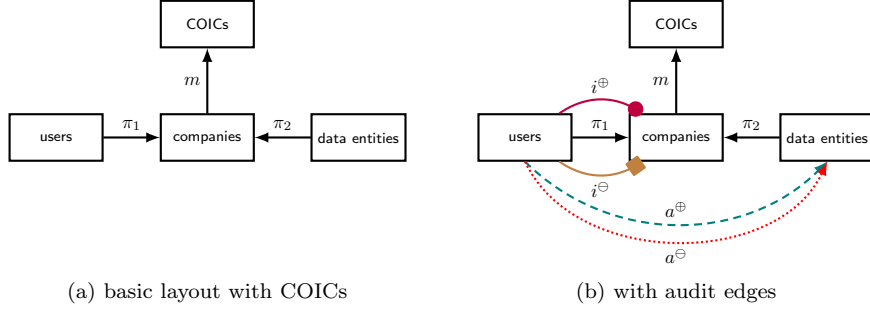


Figure 3: Chinese Wall Generalisation

To support the Chinese Wall constraint, the basic layout is supplemented by conflict of interest classes, to which companies are connected directly by the member (m) relationship (see Figure 3a). We assume here that membership of conflict of interest classes is determined when the system graph is initially populated and remains fixed through the lifetime of the system. When users are authorized (or denied) access to particular data entities, authorized (or denied) decision audit edges will result for these requests as shown in Figure 3b. We additionally introduce interest audit edges into the system graph which are added between users and companies (see Figure 3b). Active interest audit edges are labelled with i^\oplus , blocked interest audit edges are labelled with i^\ominus . We, therefore, extend the set of relationships to include the set $\{i^\oplus, i^\ominus\}$, thus allowing the system graph to support these new edges. Graphically, we represent active interest audit edges with a filled circle head, whilst blocked interest audit edges have a filled square head.

Informally, when a subject's request to access a company's data is authorized, an active interest audit edge is added (if it doesn't already exist) between the subject and the company whose data was accessed. (We will also add an authorized decision audit edge between the subject and the data entity if it does not already exist.) Additionally, blocked interest audit edges are added (if they don't already exist) between the subject and all other companies who

²It should be noted that Figure 3 does not show system graphs; it shows high-level representations of the 'shape' of a system graph.

are members of the conflict of interest class to which the first company is a member. Interest audit edges are not added after denied authorization requests.

For a concrete example, consider the system graph G_4 shown in Figure 4a, where a member of staff u_1 works for an employer e_1 . This employer supplies numerous clients (c_1 , c_2 and c_3) which have data in the form of files (f_1 , f_2 , f_3 and f_4). In this example users are connected to companies by $\pi_1 = w ; s$ whilst data entities are connected to companies by $\pi_2 = d$.

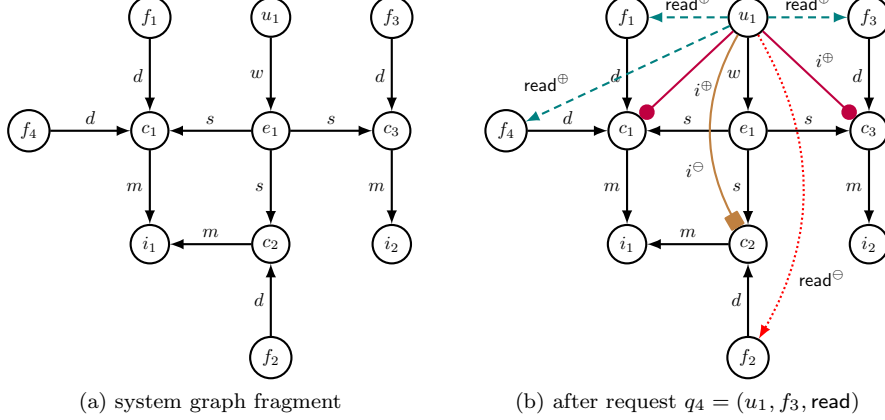


Figure 4: Enforcing the Chinese Wall policy in RPPM

If we assume the existence of a principal-matching policy $\rho = [(w ; s ; \bar{d}, p)]$ and authorization policy $PA = [(p, \star, \text{read}, 1)]$, then u_1 would be authorized to read all files. However, the clients are members of conflict of interest classes (i_1 and i_2) with clients c_1 and c_2 being competitors in i_1 . Accordingly, we modify the principal-matching and authorization policies as follows:

$$\rho_{cw} = [(i^\ominus ; \bar{d}, p_{cw}), (w ; s ; \bar{d}, p)] \quad \text{and} \quad PA_{cw} = [(p_{cw}, \star, \star, 0), (p, \star, \text{read}, 1)].$$

We now consider four different types of request that can arise. Figure 4b shows the graph G_4 after all four requests have been made.

Initial declaration of interest. The request $q_1 = (u_1, f_1, \text{read})$ to read data belonging to client c_1 in the graph G_4 will be authorized: the first principal-matching rule is not matched but the second one is. Thus $MP = [p]$ and the request is authorized, resulting in an authorized decision audit edge $(u_1, f_1, \text{read}^\oplus)$ being added to the graph along with the interest edges (u_1, c_1, i^\oplus) and (u_1, c_2, i^\ominus) .

Continued interest. If u_1 makes a second request $q_2 = (u_1, f_4, \text{read})$ for data of client c_1 this will also be authorized. The first principal-matching rule cannot be matched, as before, and the second can with principal p , once again, being authorized to read all objects. The authorized decision audit edge $(u_1, f_4, \text{read}^\oplus)$ will be added to the graph but no new interest edges are added as the required edges already exist.

Conflict of interest request. If u_1 requests data for a competing client c_2 using a third request $q_3 = (u_1, f_2, \text{read})$, this will be denied. This time, the principal-matching rule $(i^\ominus ; \bar{d}, p_{cw})$ is matched and p_{cw} is denied all actions on all objects. A denial audit edge $(u_1, f_2, \text{read}^\ominus)$ is added to the graph.

New declaration of interest which doesn't conflict. Lastly u_1 makes a request $q_4 = (u_1, f_3, \text{read})$ for data of a third client c_3 who does not conflict with c_1 (with membership in a different conflict of interest class). As with the first two requests, this request will be authorized using the second principal-matching rule and matched principal p . An authorized decision audit edge $(u_1, f_3, \text{read}^\oplus)$ is then added to the graph along with the active interest edge (u_1, c_3, i^\oplus) . No blocked interest edges are added as there are no companies other than c_3 who are members of conflict of interest class i_2 .

More generally, suppose we have a principal-matching policy ρ . In order to enforce the Chinese Wall constraint using this basic layout we add a new principal-matching rule $(i^\ominus; \overline{\pi_2}, p_{cw})$ to ρ to produce a new policy ρ_{cw} . The p_{cw} principal is denied all actions on all data entities through the inclusion of an authorization rule $(p_{cw}, \star, \star, 0)$ into the existing authorization policy PA , producing a new authorization policy PA_{cw} .³

Proposition 2. *Given an RPPM Chinese Wall constraint, as described above, for any user u the request (u, o, a) is allowed if the request is authorized by ρ_{cw} and PA_{cw} and the user u does not have an active interest in any company c' which is a member of the same conflict of interest class as the company $c \neq c'$ responsible for o . In all other cases the request is denied.*

Proof. The proof proceeds by induction on the number of evaluated requests. Consider the (base) case when no requests have yet been made. A request $q_1 = (u', o, a_j)$ will not match the inserted principal-matching rule $(i^\ominus; \overline{\pi_2}, p_{cw})$ as no blocked interest audit edges currently exist in the system graph. By assumption, request q_1 will match a preexisting principal-matching rule with principal-matching condition $\pi_1; \overline{\pi_2}$ and in doing so will match principal p . Also by assumption, principal p is authorized to perform action a_j , therefore, the audit edge (u', o, a_j^\oplus) will be added to the system graph. Additionally, the active interest edge (u', c, i^\oplus) will be added where c represents the company the target data entity o belongs to (i.e. there is a path of relations satisfying π_2 between o and c as required by the basic layout). Lastly, blocked interest edges (u', c', i^\ominus) will be added for each company $c' \neq c$ in the same conflict of interest class; these companies are identified through the existence of edges (c', i', m) in the system graph where there is also an edge (c, i, m) with $i' = i$.

Now consider the case when request $q_{x+1} = (u'', o', a_k)$ is made after x requests have been successfully evaluated. We assume, without loss of generality, that data entity o' belongs to company c_1 , a member of conflict of interest class i_1 .

- If user u'' has no active interests in any company, then request q_{x+1} will not match the inserted principal-matching rule $(i^\ominus; \overline{\pi_2}, p_{cw})$ as no blocked interest audit edges currently exist in the system graph for u'' . By assumption, request q_1 will match a preexisting principal-matching rule with principal-matching condition $\pi_1; \overline{\pi_2}$ and in doing so will match principal p . Also by assumption, principal p is authorized to perform action a_k , therefore, the audit edge (u'', o', a_k^\oplus) will be added to the system graph. The active interest edge (u'', c_1, i^\oplus) will be added to the system graph, as will blocked interest edges (u'', c_y, i^\ominus) for each company $c_y \neq c_1$ who is a member of the conflict of interest class i_1 .
- If user u'' has an active interest in company c_1 , then request q_{x+1} will not match the inserted principal-matching rule $(i^\ominus; \overline{\pi_2}, p_{cw})$ as an active, rather than blocked, interest audit edge exists between u'' and c_1 . By assumption, request q_1 will match a preexisting principal-matching rule with principal-matching condition $\pi_1; \overline{\pi_2}$ and in doing so will match

³Once again, whilst we use the AllMatch PMS and the DenyOverride CRS we could equally employ the First-Match PMS with any CRS as long as we ensure that the constraint principal-matching rules are added before any existing rules.

principal p . Also by assumption, principal p is authorized to perform action a_k , therefore, the audit edge (u', o', a_k^\oplus) will be added to the system graph. The active interest edge (u'', c_1, i^\oplus) will not be added to the system graph as it already exists. Blocked interest edges (u'', c_y, i^\ominus) for each company $c_y \neq c_1$ who is a member of the conflict of interest class i_1 will be added where they do not already exist.

- If user u'' has an active interest in company c_2 which is a member of the same conflict of interest class i_1 as c_1 , then request q_{x+1} will, in this instance, match the inserted principal-matching rule $(i^\ominus; \overline{\pi_2}, p_{cw})$. As the principal p_{cw} applies the inserted authorization rule $(p_{cw}, \star, \star, 0)$ overrides the assumed authorization achieved through principal p . The denied decision audit edge (u', o', a_k^\ominus) will be added to the system graph and no interest audit edges will be added.
- If user u'' has an active interest in company c_3 which is a member of a different conflict of interest class i_2 to c_1 but no active interest in any company which is a member of the conflict of interest class i_1 to which c_1 is a member, then request q_{x+1} will not match the inserted principal-matching rule $(i^\ominus; \overline{\pi_2}, p_{cw})$ as no interest audit edge exists between u'' and c_1 . By assumption, request q_1 will match a preexisting principal-matching rule with principal-matching condition $\pi_1; \overline{\pi_2}$ and in doing so will match principal p . Also by assumption, principal p is authorized to perform action a_k , therefore, the audit edge (u', o', a_k^\oplus) will be added to the system graph. The active interest edge (u'', c_1, i^\oplus) will be added to the system graph, as will blocked interest edges (u'', c_y, i^\ominus) for each company $c_y \neq c_1$ who is a member of the conflict of interest class i_1 .

□

The basic model described above is consistent with that used by Brewer and Nash, where there is a simple and fixed relationship between users and companies (path condition π_1) and between data objects and companies (path condition π_2). However, this approach is unnecessarily restrictive (and was chosen for ease of exposition), in that we may wish to define more complex authorization requirements between such entities. In practice, there is no reason why multiple path conditions cannot be used between users, objects and companies, each of which is mapped to the appropriate principal.

For example, given two paths of relations between users and companies ($w; s$ and $w; p; s$) and two paths of relations between data entities and companies (d and $f; d$) the principal-matching policy from our running example is modified to include both options for blocking paths and all combinations for normal authorization.

$$\begin{aligned} \rho_{cw_2} = & [(i^\ominus; \overline{d}, p_{cw_2}), (i^\ominus; \overline{d}; \overline{f}, p_{cw_2}), \\ & (w; s; \overline{d}, p), (w; s; \overline{d}; \overline{f}, p), (w; p; s; \overline{d}, p), (w; p; s; \overline{d}; \overline{f}, p)] \end{aligned}$$

5 Related Work

Relationship-based access control is becoming an increasingly important alternative approach to specifying and enforcing authorization policies. A number of models have been proposed in recent years [4, 5, 6, 7, 9, 17], but most have focused on access control in social networks [4, 5, 6, 9, 17]. In this paper, we extend the RPPM model of Crampton and Sellwood [7] by introducing additional types of edges to support efficient request evaluation and history-based access control policies.

History-based access control, where an authorization decision is dependent (in part) on the outcome of previous requests, has been widely studied since Brewer and Nash's seminal paper

on the Chinese Wall policy [3]. The enforcement mechanism for this policy is based on a history matrix, which records what requests have previously been allowed. It is very natural to record such information as audit edges in the system graph and to use these edges to define and enforce history-based policies. Fong *et al.* recently proposed a relationship-based model that incorporated temporal operators, enabling them to specify and enforce history-based policies [10]. This work extended Fong’s ReBAC model, developed in the context of social networks, and is thus unsuitable for the more generic access control applications for which the RPPM model was designed. In particular, there is no obvious way in which it can support Chinese wall policies.

There has been some interest in recent years in reusing, recycling or caching authorization decisions at policy enforcement points in order to avoid recomputing decisions [2, 12, 13, 16]. In principle, these techniques are particularly valuable in large-scale, distributed systems, providing faster decisions, and the potential to recover from failures in the communication infrastructure or failure of one of the components, in the (distributed) authorization system. However, many of the techniques are of limited value because the correlation between access control decisions and the structure of access control policies is typically rather low. In contrast, caching in the RPPM model has the potential to substantially speed up decision-making because a cached edge is of real value as it enables the decision-making apparatus to sidestep the expensive step of principal matching and proceed directly to evaluating the authorization policy. Moreover, a cached edge applies to multiple requests, irrespective of whether the request has previously been evaluated, unlike many, if not all, proposals in the literature.

6 Conclusion

The RPPM model fuses ideas from relationship-based access control (by using a labelled graph), role-based access control (in its use of principals to simplify policy specification) and Unix (by mapping a user-object relationship to a principal before determining whether a request is authorized). This unique blend of features make it suitable for large-scale applications in which the relationships between users are a crucial factor in specifying authorization rules.

In addition to these advantages, the RPPM model is particularly suitable for recording information that may be generated in the process of making authorization decisions. In this paper, we focus on two new types of edges. Caching edges introduce shortcuts in the system graph indicating the principals associated with a user-object pair. Such edges can introduce substantial efficiencies to the evaluation of decisions. Audit edges allow for the enforcement of history-based policies, including separation of duty and Chinese wall policies.

The introduction of audit edges lays the foundation for future work supporting workflow tasks within the RPPM model. This work may, additionally, require the model to be further extended with the introduction of stateful entities.

References

- [1] M. Abadi and C. Fournet. Access control based on execution history. In *NDSS*. The Internet Society, 2003.
- [2] K. Borders, X. Zhao, and A. Prakash. CPOL: high-performance policy evaluation. In V. Atluri, C. Meadows, and A. Juels, editors, *ACM Conference on Computer and Communications Security*, pages 147–157. ACM, 2005.
- [3] D. F. C. Brewer and M. J. Nash. The Chinese Wall security policy. In *IEEE Symposium on Security and Privacy*, pages 206–214. IEEE Computer Society, 1989.

- [4] B. Carminati, E. Ferrari, and A. Perego. Enforcing access control in web-based social networks. *ACM Trans. Inf. Syst. Secur.*, 13(1), 2009.
- [5] Y. Cheng, J. Park, and R. S. Sandhu. Relationship-based access control for online social networks: Beyond user-to-user relationships. In *SocialCom/PASSAT*, pages 646–655. IEEE, 2012.
- [6] Y. Cheng, J. Park, and R. S. Sandhu. A user-to-user relationship-based access control model for online social networks. In N. Cuppens-Boulahia, F. Cuppens, and J. García-Alfaro, editors, *DBSec*, volume 7371 of *Lecture Notes in Computer Science*, pages 8–24. Springer, 2012.
- [7] J. Crampton and J. Sellwood. Path conditions and principal matching: a new approach to access control. In S. L. Osborn, M. V. Tripunitara, and I. Molloy, editors, *SACMAT*, pages 187–198. ACM, 2014.
- [8] G. Edjlali, A. Acharya, and V. Chaudhary. History-based access control for mobile code. In J. Vitek and C. D. Jensen, editors, *Secure Internet Programming*, volume 1603 of *Lecture Notes in Computer Science*, pages 413–431. Springer, 1999.
- [9] P. W. L. Fong. Relationship-based access control: protection model and policy language. In R. S. Sandhu and E. Bertino, editors, *CODASPY*, pages 191–202. ACM, 2011.
- [10] P. W. L. Fong, P. Mehregan, and R. Krishnan. Relational abstraction in community-based secure collaboration. In A.-R. Sadeghi, V. D. Gligor, and M. Yung, editors, *ACM Conference on Computer and Communications Security*, pages 585–598. ACM, 2013.
- [11] V. D. Gligor, S. I. Gavrilă, and D. F. Ferraiolo. On the formal definition of separation-of-duty policies and their composition. In *IEEE Symposium on Security and Privacy*, pages 172–183. IEEE Computer Society, 1998.
- [12] M. Kohler, A. D. Brucker, and A. Schaad. Proactive caching: Generating caching heuristics for business process environments. In *CSE (3)*, pages 297–304. IEEE Computer Society, 2009.
- [13] M. Kohler and R. Fies. Proactive caching - a framework for performance optimized access control evaluations. In *POLICY*, pages 92–94. IEEE Computer Society, 2009.
- [14] K. Krukow, M. Nielsen, and V. Sassone. A logical framework for history-based access control and reputation systems. *Journal of Computer Security*, 16(1):63–101, 2008.
- [15] R. T. Simon and M. E. Zurko. Separation of duty in role-based environments. In *CSFW*, pages 183–194. IEEE Computer Society, 1997.
- [16] Q. Wei, J. Crampton, K. Beznosov, and M. Ripeanu. Authorization recycling in hierarchical rbac systems. *ACM Trans. Inf. Syst. Secur.*, 14(1):3, 2011.
- [17] R. Zhang, A. Artale, F. Giunchiglia, and B. Crispo. Using description logics in relation based access control. In B. C. Grau, I. Horrocks, B. Motik, and U. Sattler, editors, *Description Logics*, volume 477 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2009.